

This application is submitted in the name of inventors Sunil Chandrupatla, Aravind Sitaraman and Leslie Thomas, assignors to Cisco Technology, Inc., a California Corporation.

SPECIFICATION

AGGREGATION OF USER USAGE DATA FOR ACCOUNTING SYSTEMS IN DYNAMICALLY CONFIGURED NETWORKS

BACKGROUND OF THE INVENTION

1. Field Of The Invention

The present invention relates to an account metering method and apparatus for use in a computer network. More particularly, the present invention relates to a method and apparatus for obtaining and correlating account metering data from two distinct sources: start-stop event accounting records associated with accounting servers and detailed flow data collected from numerous routers throughout the network environment.

2. Background

The ability to provide computer networking capabilities to the home personal computer (PC) is most commonly provided by telephone companies (Telcos) or commercial Internet Service Providers (ISPs) that maintain network operation centers (NOCs) along the information superhighway. Network operation centers, commonly located within wide area networks (WANs), serve to house the network interfaces and service components necessary to provide routing, bridging and other essential networking functions. It is via these network operation centers that the user or subscriber, through a host computer connected to an access point, can connect with

public domains, such as the Internet and private domains, such as intra-networks and community-of-interest (pay-for-use) domains.

Currently, Telcos and ISPs are limited in the manner that they can charge customers for a product. Basically, Telcos and ISPs are confined to either charging a flat fee on a monthly basis, thus allowing the user unlimited network access for the specified period, or charging the user on a rate basis, typically, an hourly rate. These billing schemes are unsophisticated because they provide for only a simplified method of accounting for the wide spectrum of events that a user undertakes while the user is logged on to the access point. Current technology only allows for the Telco or ISP to account for the duration of the period from when a user logs on to and when the user subsequently logs off.

As an example, FIG. 1 schematically illustrates a current model of a network accounting system. In this typical networking environment 10, a user 12 implements a "dashboard" application on a host/computer 14 that requires them to input identification and authorization information, such as a user name and a password. This information is then sent, via the host having a dial-up connection, to the Telco or ISP operated NOC 16. A dial-up connection access method is implemented by a modem (not shown) located at the host and a modem pool (not shown) maintained at the ISP. Those of ordinary skill in the art will recognize that other types of access methods may be provided by an ISP such as frame relay, leased lines, ATM (asynchronous transfer mode), ADSL, and the like. A network access server (NAS) 18 located at the NOC 16 receives the identification and authorization information and proxies it to an authentication,

authorization and accounting (AAA) server 20. Once the server verifies the user authentication and authorization it grants the user log-on access to public domains or private domains, such as the Internet 22. At the moment authorization is granted, counters 24 within the NAS 18 are engaged which begin tracking the duration of the session as well as the byte count encountered during the session. The record of this account log-on event is then forwarded to the AAA server 20. Subsequently, when the user desires to log-off or a log-off is warranted by other factors outside the control of the user, the counters 24 within the NAS 18 are stopped and the appropriate accounting log-off event is forwarded to the AAA server 20. Thus, the accounting information provided to the service provider is limited to a subscriber's account log-on and log-off activities (start and stop times) and byte counts. This severely limits the billing schemes that may be implemented.

The Telco or ISP, as well as the subscribers, would benefit from having a more developed account metering scheme that makes possible alternate billing options. The service provider, who has real-time access to such detailed subscriber account activity, can then devise and offer customized billing rates and level-of-service schemes to subscribers based on a variety of factors. For example, billing schemes could be devised based upon what services a user accesses (i.e. Internet versus private domain services) and the duration of such connections (which might differ from the overall connection duration). Billing could also be provided based upon the priority given to the data transmitted (i.e. text versus video or audio) or the byte count intensity of specific flows. Additionally, service providers could charge customers based on the transmission roadmap (i.e. number of hops encountered for a given flow). The potential for extensive

and diverse rating schemes and billing schemes becomes apparent if the service provider has efficient access to comprehensive accounting records. The subscriber benefits because, instead of the service provider spreading costs equally among the entire subscriber base, rates can be modified so that the subscriber who makes less costly connections can be charged commensurately less. As networking systems continue to scale up and the number of service providers grow, competition and subscriber demand will put pressure on providers to implement billing rates and service plans that are more tailored to the individual requisites of the users. Network service providers confronted with this problem will want systems that overcome these issues and provide detailed subscriber accounting information in an automated concise format.

As networking systems continue to scale-up, it will also become vital for the service providers to have efficient access to accounting records as a way of assessing user trends and adjusting the network configuration in accordance with these trends. As the information superhighway becomes increasingly more congested, the ability of the service providers to facilitate network traffic becomes a heightened concern. If the service providers have the necessary information readily available to foresee user trends and bottlenecks in transmission, then appropriate modifications can be made; for example, hardware may be added, software may be modified or traffic may be rerouted. Ultimately, the subscribers benefit by receiving a higher quality and more reliable service provided to them.

Providing account metering on Packet Switched Networks (PSNs) is complicated because no integrated accounting mechanism exists to resolve information originating

from multiple sources throughout the PSN structure. The ISPs and Telcos would benefit from an account metering system capable of aggregating accounting-related data originating from various sources on the composite PSN, capturing the data efficiently and correlating the data to provide a comprehensive subscriber-specific accounting record. For instance, network flow data can be collected at downstream routers located at Points of Presence (PoPs). This data contains detailed traffic statistics, such as a timestamp of a specific flow, source and destination IP addresses, source and destination port numbers, next hop addresses, total byte count in a specific flow, and the like. However, since the source of such flow data originates from numerous PoPs, current technology has not provided for this data to be collected, filtered and aggregated efficiently so that it can benefit service providers in devising detailed account rating and billing schemes. Additionally, this flow data has no accounting benefit unless it can be matched against the user-specific accounting start-stop events recorded at the accounting servers. Therefore, the need exists to combine efficiently the accounting event data with the network flow data to result in one unified comprehensive accounting record that details both user start-stop accounting events and the flow information corresponding to specific user sessions.

BRIEF DESCRIPTION OF THE INVENTION

A method and apparatus for providing an aggregated account metering system to a computer network service provider resulting in comprehensive detailed subscriber accounting records. Accounting start-stop event data is retrieved from accounting servers. The accounting records are forwarded to a first adapter where they are parsed and then published on an active information bus. Network flow data is retrieved from

network flow collectors that are in communication with routers. The network flow data is forwarded to a second adapter where it is aggregated and then published on an active information bus. An integrating accounting adapter subscribes to and collects accounting event data and network flow data and correlates this data into a detailed call record formatted as desired.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a schematic drawing of a computer network illustrating a prior art approach to log-on and log-off account metering in accordance with the prior art.

FIG. 2 is a schematic drawing of a computer network illustrating the real-time conveyance and aggregation of accounting related data from both accounting servers and network flow collectors in accordance with a presently preferred embodiment of the present invention.

FIG. 3 is a flow diagram of a method for real-time conveyance and aggregation of accounting related data in a packet switch network in accordance with a presently preferred embodiment of the present invention.

DETAILED DESCRIPTION OF THE PRESENT INVENTION

Those of ordinary skill in the art will realize that the following description of the

present invention is illustrative only and is not intended to be in any way limiting. Other embodiments of the invention will readily suggest themselves to such skilled persons from an examination of the within disclosure.

FIG. 2 illustrates schematically a dynamically configured packet switch network (PSN) 50 embodying the aggregated accounting system in accordance with a presently preferred embodiment of the present invention. Accounting related information originates from two distinct sources within the PSN 50. One source is accounting servers that store subscriber start-stop events. The other is from flow collectors that store network flow usage records. The present invention retrieves, aggregates and publishes these records to a centralized accounting interface where they are synthesized to form a unified detailed call record.

A subscriber 52 by way of a computer/host 54 is granted access to a packet switch network 50 through network service providers. Network service providers, such as Internet Service Providers (ISPs) or Telephone Companies (Telcos), maintain Points of Presence (PoPs) 58, 60 and 62 and network operation centers (NOC) 56. PoPs 58, 60 and 62 are geographical areas serviced by NOCs which are managed by an ISP or a Telco. Located within the PoPs 58, 60 and 62 are network access servers (NASs) 64, 66 and 68 that grant network access to subscribers 52 through authorization and authentication processes. The NASs 64, 66 and 68 are in communication with accounting servers 70, 72 and 74 which may be physically located at the same location as the corresponding NAS, as shown in FIG. 2, or the accounting servers may be located at the NOC or elsewhere. The location and the quantity of accounting servers are

dictated by a number of factors including the complexity of the system, the geographical span of the service provider's system and network bandwidth concerns.

For example, the accounting servers may include an authentication, authorization and accounting (AAA) server, such as Cisco ACS or Cisco Secure, manufactured by Cisco Systems, Inc. of San Jose, CA or any other similar standard accounting server. In connection with the accounting servers 70, 72 and 74 are storage devices 76, 78 and 80 that serve as memory storage devices for the flat file accounting records proxied from the accounting servers 70, 72 and 74. The NASs 64, 66, and 68 communicate with the accounting servers 70, 72 and 74 by means of a standardized internet protocol such as, for example, the Remote Authentication Dial-In User Service (RADIUS) protocol. Other protocols, known by those of ordinary skill in the art, can also be used to communicate accounting start/stop event data between the NASs 64, 66 and 68 and accounting servers 70, 72 and 74.

Once the subscriber 52 has been granted network access through a successful authorization and authentication process, for example, at PoP 58, the NAS 64 or another interface located at PoP 58, will simultaneously generate accounting start packets and forward such to an accounting server 70. It is feasible, and within the inventive concept herein disclosed, to generate and forward such start/stop event packets from within other hardware located within the PoP 58. For example, accounting start/stop event requests may be generated and forwarded by a gateway device, such as Cisco 6510 Service Selection Gateway (SSG) manufactured by Cisco Systems, Inc. of San Jose, CA. In accordance with a presently preferred embodiment of the present invention the

RADIUS account start packet will have the following attributes associated with the record:

Acct-Status-Type = Start
 NAS-IP-Address = *ip_address*
 User-Name = "username"
 Acct-Session-Id = "session_id"
 Framed-IP-Address = *user_ip*
 Proxy-State = "n"

where:

ip_address = IP address of the SSG interface card 1.
username = Name used to log on to the service provider network
session_id = Session Number
user_ip = IP address of the user's system
n = Accounting record queuing information

Once the subscriber/user 52 completes a log-off process or other factors outside the subscriber's control cause a log-off to occur, the NAS 64, or another interface located at the PoP 58, will generate an account stop packet and forward such to an accounting server 70. In accordance with a presently preferred embodiment of the present invention the RADIUS accounting stop packet will have the following attributes associated with the record:

09198500 113393
 060506T50

Acct-Status-Type = Stop
 NAS-IP-Address = ip_address
 User-Name = "username"
 Acct-Session-Time = time
 Acct-Terminate-Cause = cause
 Acct-Session-Id = "session_id"
 Framed-IP-Address = user_ip
 Proxy-State = "n"

where:

ip_address = IP address of the SSG interface card 1.
username = Name used to log on to the service provider network.
time = Length of session in seconds.
cause = Cause of account termination. These include:
 - User-Request - Session-Timeout
session_id = Session Number.
user_ip = IP address of the user's system.
n = Accounting record queuing information.

The accounting start-stop event data is not limited to account log-on and account log-off events. Other start-stop events that a subscriber can trigger are also feasible and within the inventive concept herein disclosed. These events include, but are not limited to; (1) a subscriber establishing or terminating a tunnel-based service such as an L2F (Layer Two Forwarding) or L2TP (Layer Two Tunneling Protocol) service, for example; and (2) a subscriber establishing or terminating a Point to Point Protocol (PPP) connection. Interfaces at the PoP 58 will generate the accounting start and stop packets for each of these events and forward the packets to the accounting server 70. The accounting start-stop packets are capable of accounting for the duration of these events, as well as the number of bytes encountered for any of these events.

As data accumulates within the storage device 76 associated with corresponding accounting server 70, the accounting adapter 82 will periodically parse the accounting

start-stop event data from the storage device 76. The accounting adapter 82 will be configured to parse accounting records from the storage device 76 at prescribed time intervals, most commonly defined in terms of seconds. The accounting records are then published as accounting event data by the accounting adapter 82 on to an information bus 88. The information bus 88 can be Common Object Request Broker Architecture (CORBA)-based, which handles the communication of messages to and from objects in a distributed, multi-platform environment, or another acceptable communication means can be used as known by those of ordinary skill in the art.

CORBA provides a standard way of executing program modules in a distributed environment. The broker, therefore, may be incorporated into an Object Request Broker (ORB) within a CORBA compliant network. To make a request of an ORB, a client may use a dynamic invocation interface (which is a standard interface which is independent of the target object's interface) or an Object Management Group Interface Definition Language (OMG IDL) stub (the specific stub depending on the interface of the target object). For some functions, the client may also directly interact with the ORB. The object is then invoked. When an invocation occurs, the ORB core arranges so a call is made to the appropriate method of the implementation. A parameter to that method specifies the object being invoked, which the method can use to locate the data for the object. When the method is complete, it returns, causing output parameters or exception results to be transmitted back to the client.

For a typical accounting event published on the information bus 88 the following data fields will be present:

Table 1

Field	Type	Comments
strNasId	string	The NAS address where the subscriber is logging.
strTimeStamp	string	Time when the packet is published.
guidSource	string	GUID of the accounting adapter.
guidTarget	string	GUID of the target device subscribing to the accounting adapter.
accountingPacket	structure of AV pairs	Includes structure of AV pairs; name and value.

A second source of accounting-related data is generated throughout the PSN 50 at the routers 90, 92 and 94 located within the PoPs 58, 60 and 62. Routers located throughout the PSN 50 serve to forward and direct network traffic between the users and the domains. The data packets that flow through the routers 90, 92 and 94 contain detailed flow data such as Layer 2, Layer 3 and Layer 4 flow information contained in respective packet headers. The routers are configured to send this detailed flow data to network flow collectors 96, 98 and 100. Numerous routers may exist in any one given PoP, so that the network flow collectors are in communication with any number of routers within the PoP. The network flow collectors 96, 98 and 100 collect the flow data and log the information onto flat files at storage devices 102, 104 and 106.

The network flow data exported by a network flow collector comprises expired traffic flows that contain specific traffic statistics. These traffic flows contain detailed

information about network layer sources and destinations, including the level of individual applications and protocols constituting the end-to-end conversation. For example, the fields included in the detailed traffic statistics may include the following: timestamp of the flow, source and destination IP addresses, source and destination port numbers, input and output interface number, next hop address, number of packets in the flow and the first and last timestamps of packets in a particular flow. In accordance with a presently preferred embodiment of the present invention the network flow collectors use an aggregation scheme based on IP addresses. An example of the contents of a network flow usage file follows:

ROUTER 171.69.73.146/TYPE IP/AGGREGATION SourceNode/PERIOD

10/UTC_Begin 869784176/UTC_End 869784776

where:

ROUTER 171.69.73.146	=	The IP address of the router from which the network flow file was sent.
TYPE IP	=	Type of protocol used.
AGGREGATION SourceNode	=	Aggregation scheme used.
PERIOD 10	=	The time period between which the network flow collector logs information.
UTC_Begin	=	Beginning time for the collection of data.
UTC_End	=	End time for the collection of data.

Subsequent lines in the file read as follows:

172.22.6.63/103/7828/33

where:

172.22.6.63	=	Source IP address.
103	=	Sum of all packets sent.
7828	=	Total number of bytes (octets) sent from the source
33	=	Number of flows aggregated into this flow.

The comprehensive detailed traffic statistics can then be organized by

aggregation schemes or the data volume reduced by filtering schemes. For example, aggregation schemes can be devised based upon the following parameters: a specified time period, source or destination address present in the flow export data, source or destination ports, matching source and destination pairs in the input data, protocols present in the input data, interfaces present in the input data or every communicating pair in the input data. Filtration can permit or deny the further processing of network flow usage data as a way of customizing the network flow usage data. For example, the following data fields can be filtered: Layer 2 MAC source or destination address, Layer 3 IP source or destination address, Layer 4 input or output interface numbers, Layer 4 source or destination port, and the like.

The network flow data is parsed from the a network flow collector storage device 102 by network flow adapter 108. Identification information within the network flow file will allow for the network flow adapter 108 to parse only the network flow data that it subscribes to. The network flow adapter 108 will be configured to parse network flow data from the network flow collector storage device 102 at prescribed time intervals, most commonly defined in terms of seconds. Once the network flow adapter 108 parses the files they create network flow events and publish the events on an information bus 88. The information bus can be CORBA-based, as discussed above, or another acceptable communication means. For a typical network flow event published on the information bus the following data fields will be present:

Table 2:

Field	Type	Comments
routerName information	string	IP address of router exported.
flowTime	integer	hhmm(hour/minute) when file is logged.
guidSource	string	GUID of the Network Flow Adapter.
guidTarget	string	GUID of the target device subscribing to the network flow event.
flowHeader	string	First line in the Network Flow data.
flowPacket	structure of field arrays	Array contains srcIPAddress, sumPackets, totalBytes, numFlows.

Once the accounting events and the network flow usage events are published on the information bus 88 the subscribing target device 114 collects the data. The target device 114 is maintained by an ISP or Telco and is located at a NOC 56 within the PSN 50. The target device 114 will capture only those accounting events and network flow usage events that it subscribes to. An Integrated Accounting Adapter 116 incorporated with the target device 114 correlates the accounting start-stop event data and the network flow data resulting in a call detail record. Typically, the correlation accomplished within the Integrated Accounting Adapter will be to match user accounting start-stop events with corresponding user network flow data resulting in specific call detail records associated with specific users/subscribers. This type of correlation is beneficial to later rating and billing applications. Other modes of correlation can be performed that would suit the particular needs of the service provider.

For instance, correlation schemes can be used to support network planning and monitor network resources. The correlated call detail records are then forward to a storage device 118 where they are written into flat files for post processing. In most instances, post-processing applications 120 will involve a service provider maintained and unique billing application that reformats the call detail records to meet specific needs. Other post-processing applications, such as traffic analysis tools are also feasible and within the inventive concepts herein disclosed. An example of an alternate post-processing application would be an account that details at what interfaces within the network flow congestion exists. Such an application would be beneficial to the ISP or Telco in determining how flows could be re-routed or what additional hardware is needed to alleviate the problem. Those of ordinary skill in the art will recognize that the data may be organized in many other useful ways without departing from the inventive concept herein.

FIG. 3 illustrates a flow diagram of a method for implementing an aggregated account metering method in a packet switched network (PSN) in accordance with a presently preferred embodiment of the present invention. Accounting related information originates from two distinct sources within a PSN; the accounting start-stop event data stored on accounting servers and the network flow data captured at flow collectors from routers located throughout the network at Points of Presence (PoPs).

In FIG. 3 at reference number 130, the accounting start-stop event data are parsed by an accounting adapter from the various data banks associated with the accounting servers. The accounting start-stop event data typically include user specific

account log-on and account log-off intervals and byte counts encountered during specified account sessions. In addition, the accounting start-stop event data may include service start and service stop intervals and connection start and connection stop intervals, as well as corresponding byte counts for a specific interval. The parsing accounting adapter is in communication with each subdirectory within the accounting server's memory bank on prescribed intervals. The communication link between the accounting servers and accounting adapters may be conducted according to a suitable internet protocol, such as the Radius Authentication Dial-In User Service (RADIUS) protocol. Those of ordinary skill in the art will realize that other protocols can be used as an acceptable communication link between the various communication interfaces that encounter accounting start-stop event data. Once parsed, at reference number 140 the accounting start-stop event data are published onto an information bus. For example, the information bus may be CORBA-based, as discussed above in more detail, or another acceptable communication means can be used as known by those of ordinary skill in the art.

Concurrently with the accounting adapter parsing and publishing accounting start-stop event data, at reference number 150 a network flow collector compiles, via subscription, network flow data accumulated from numerous routers throughout the PoPs. At reference number 160, the network flow collector aggregates the network flow data according to a configured aggregation scheme. While the aggregation performed within the network flow collectors is not essential to the overall inventive concepts herein disclosed, aggregation serves to coordinate the data in accordance with a preferred designated parameter. A preferred aggregation scheme used in accordance

with a presently preferred embodiment of the present invention is based on IP addresses. Other aggregation schemes are also possible and within the inventive concept herein disclosed. Examples of other aggregation schemes include, but are not limited to, aggregation schemes based on the following: a specified time period, source or destination address present in the flow export data, source or destination ports, matching source and destination pairs in the input data, protocols present in the input data, interfaces present in the input data or every communicating pair in the input data. Additionally, the network flow adapter, at reference number 170, may use an optional filtering scheme that is performed for data volume reduction and further customizing the network flow data. Filtering serves to eliminate extraneous data from the further processing, lowering the byte count and making the overall aggregating account metering process more efficient. Filtering defines what data is included or excluded from further processing and multiple permit or deny filters can be imposed. Filtering can be imposed either prior to aggregation or once aggregation has been completed. By way of example, the following data fields can be filtered: network layer IP source or destination address, next hop address, input interface number, output interface number, source or destination port or protocol.

At reference number 180, the network flow adapter, at a prescribed interval, parses the various network flow collectors. The parsing act results in the network flow adapters retrieving detailed network traffic flow statistics. These traffic flows contain detailed information about network layer sources and destination, including the level of individual applications and protocols constituting the end-to-end conversation. By way of example, the fields in the detailed traffic statistics may include the following:

timestamp of the flow, source and destination IP addresses, source and destination port numbers, input and output interface number, next hop address, number of packets in the flow and the first and last timestamps of packets in a particular flow. Once the network flow data has been parsed from the flow collectors to the network flow adapters, at reference number 190 the data is published on to the information bus. As for the accounting adapter publication, the information bus may be CORBA-based or otherwise as discussed in more detail above.

At reference number 200, the subscribing target device captures the accounting start-stop event data and the network flow data from the information bus. At reference number 210, the accounting start-stop event data and the network flow data are correlated within the integrated accounting adapter. The correlation accomplished within the integrated accounting adapter will, routinely, be to match user accounting start-stop events with user flow data resulting in specific call detail records associated with specific users/subscribers. This type of correlation will be used by subsequent rating and billing applications. Other modes of correlation can be performed that would suit the particular needs of the service provider. For instance, correlation schemes can be used to support network planning and to monitor network resources. At reference number 220, the post-correlated call detail records can be optionally reformatted in a conventional manner to meet service provider's post-processing applications. Post-processing applications include service provider rating and billing applications, network planning applications and monitoring of network resources as well as other applications which can benefit from the type of information aggregated.

Although illustrative presently preferred embodiments and applications of this invention are shown and described herein, many variations and modifications are possible which remain within the concept, scope and spirit of the invention, and these variations would become clear to those skilled in the art after perusal of this application. The invention, therefore, is not limited except in spirit of the appended claims.